AMD DEVELOPER symposium 2002

Introduction to HyperTransport™ Technology

Chris Neuts
Manager of Technology Evangelism
CPG Marketing

Today's PC System Architecture

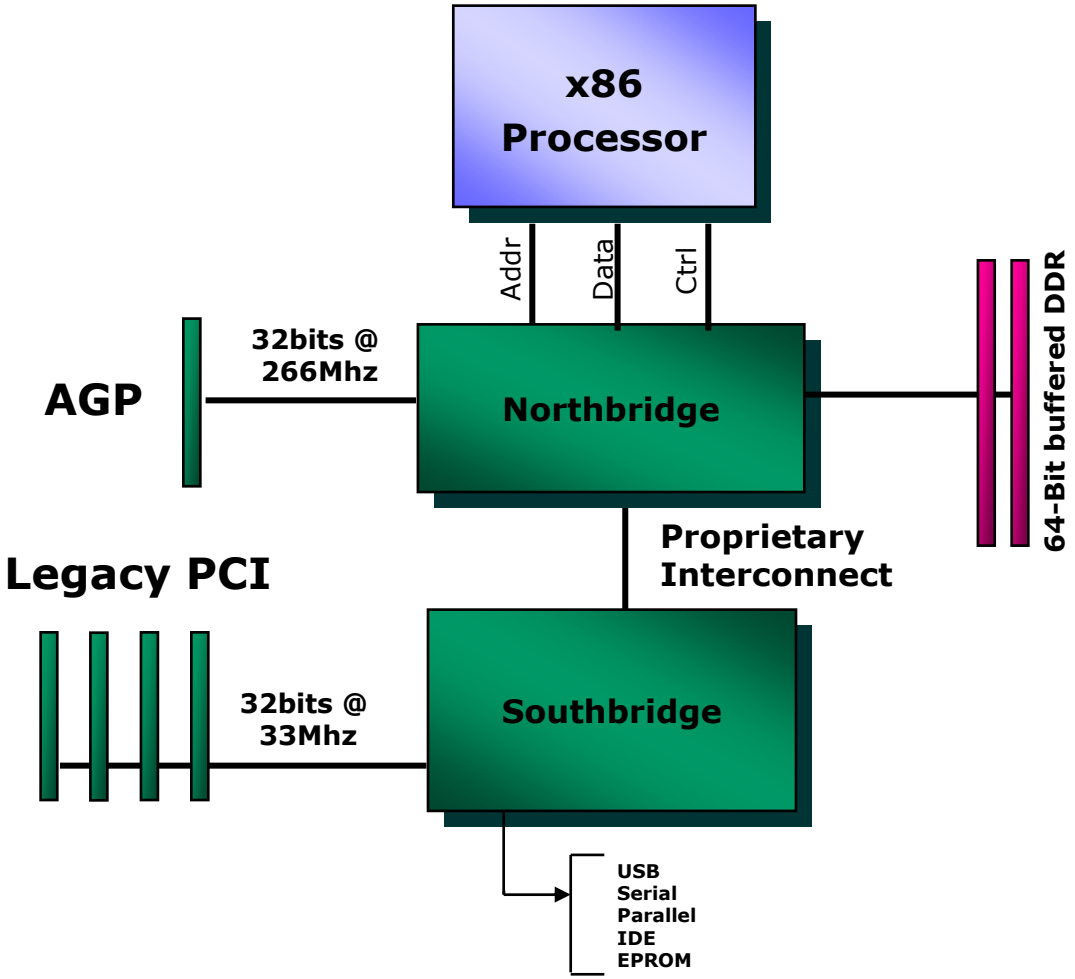

System Weaknesses

- 32 bit Processor
- Memory Latency and Bandwidth
- I/O Latency and Bandwidth
- Northbridge is the system bottleneck
- Doesn't easily support a multiprocessing architecture

The Multiprocessing Dilemma

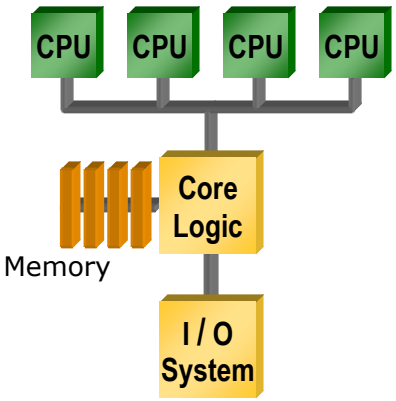

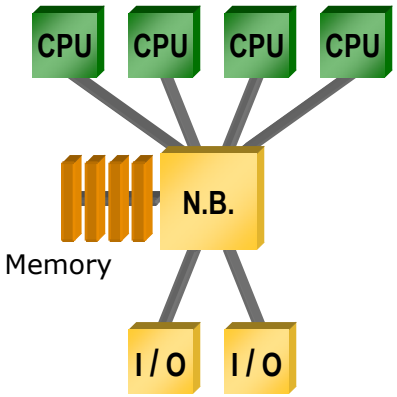

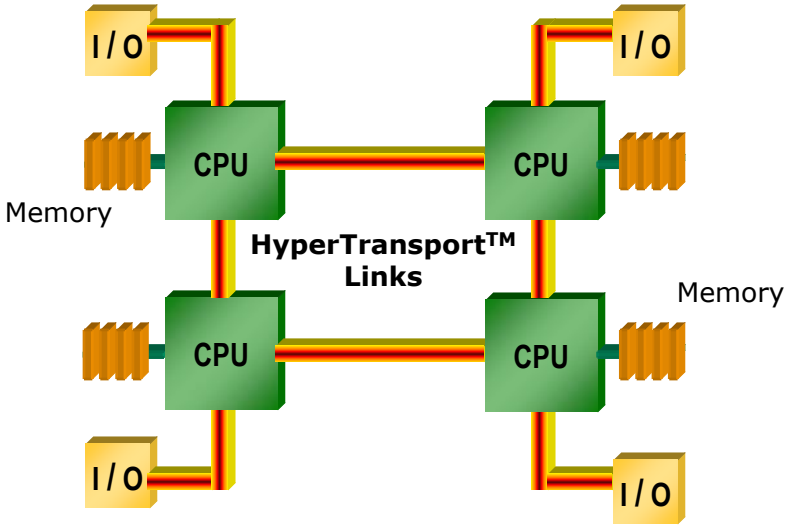

- Current server and workstation architectures face limitations
  - Shared CPU buses force multiple CPUs to share limited bandwidth
  - Adding more processors does NOT add more memory bandwidth
  - Adding more processors does NOT add more I/O bandwidth
- Point-to-point and distributed architecture improves scaling
  - Distributed architecture expands memory and I/O bandwidth as additional processors are added
  - HyperTransport™ technology supports distributed microprocessor architectures

Vision Behind HyperTransport™ Technology

AMD DEVELOPER symposium 2002

- Designed to provide excellent chip-to-chip interconnect for computation, networking, and telecommunications devices
- Scalable in speed and width to be a viable solution for battery powered devices, as well as high-end servers
- Helps to reduce I/O bottlenecks by efficiently integrating legacy buses
- Provides sufficient bandwidth for supporting new interconnects including PCI-X, InfiniBand, 10Gbit Ethernet and other evolving standards
- Offered royalty-free by the HyperTransport Technology Consortium to promote the next wave of technology innovation

Why Use HyperTransport™ Technology?

AMD DEVELOPER symposium 2002

- High performance chip-to-chip interconnect
- Scalable & flexible to cover range of performance needs and topologies
- Available today in the form of:
  - Chips – Processors and I/O ICs
  - IP (Verilog RTL)
  - Tools (Test bench and BFM) and,
  - PHY from a variety of vendors
- Layered architecture supports future protocol enhancements and higher speed physical layer
- Full PCI software compatibility
- Open Standard – no licensing fee or royalty

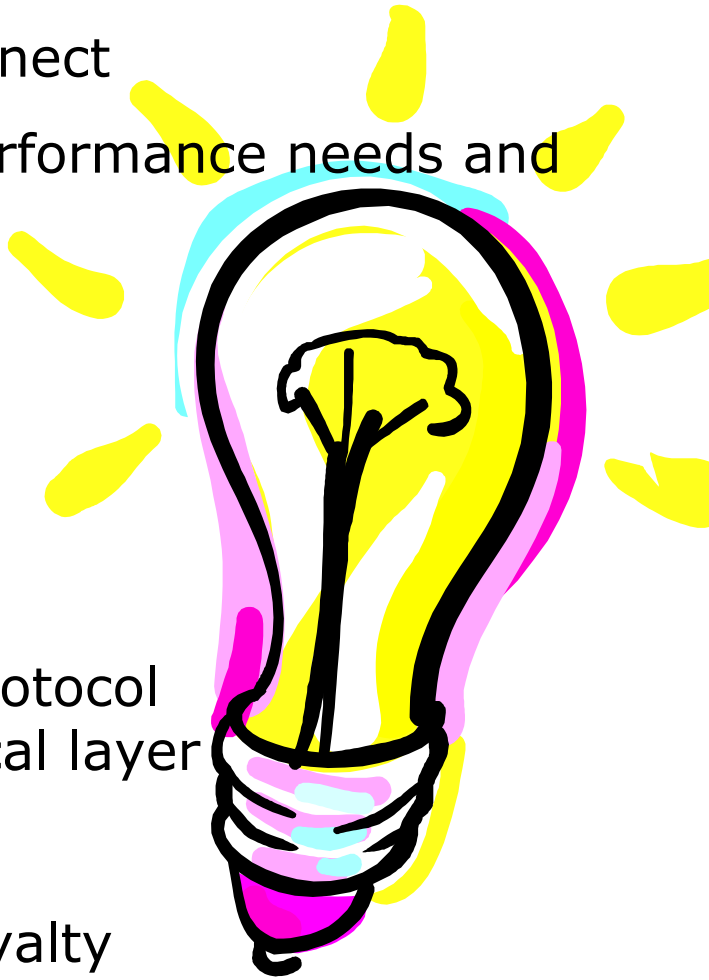

The "Virtual Chipset"

- HyperTransport™ devices can operate on different processors
- Devices can be shared between applications
  - PCs
  - Networking
  - Communications
  - Interoperable
  - Enables Large Markets and High Unit Volumes

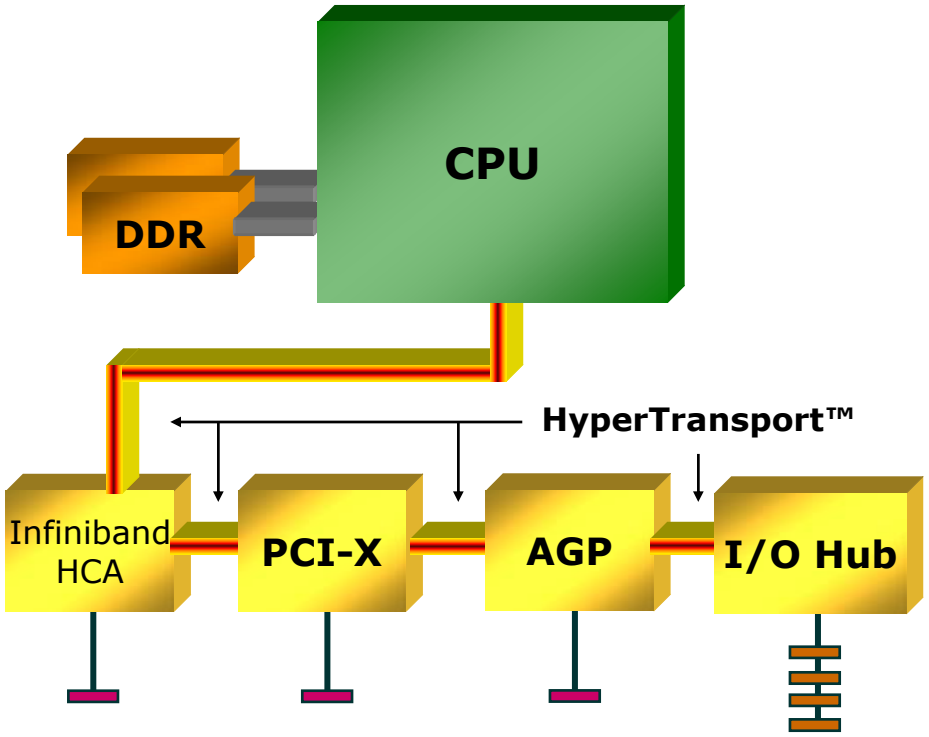

HyperTransport™ Physical Interface

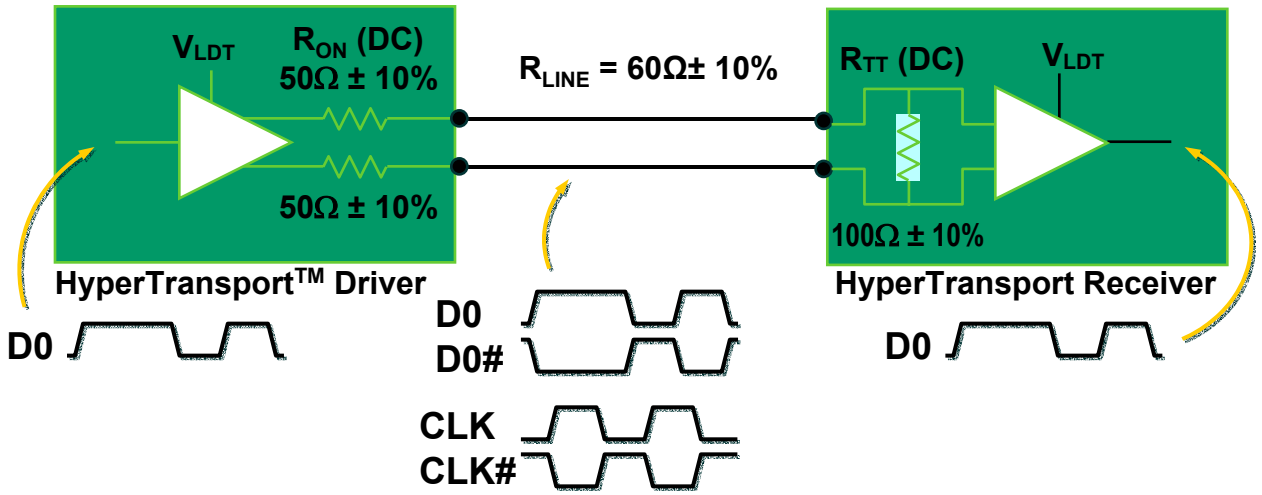

- Low voltage, differential signaling
  - Two pins per bit - pin pairs swing in opposite directions
  - VLDT is 1.2 volts ± 5% resulting in a differential output of 600 mVTYPICAL
  - Differential voltage at the receiver inputs can be as low as 200 mV
- 60 ohm differential impedance for low cost PCBs
  - No special PCB stack-up required (FR4 works)
  - Trace lengths up to 24 inches at 800MHz operation

HyperTransport™ Technology Scalable Link Width and Speed

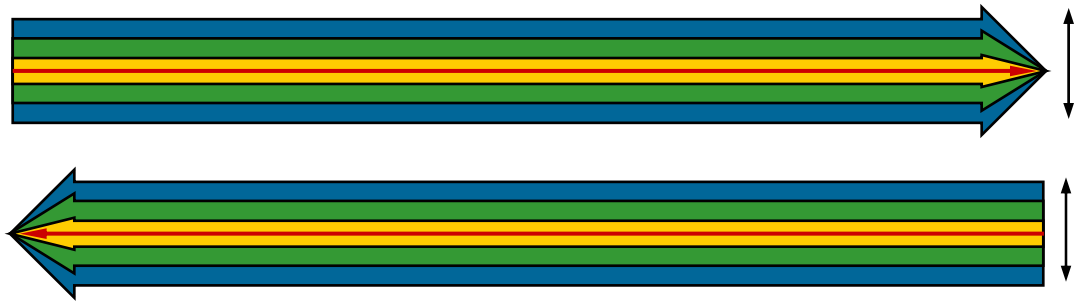

Up to 1.6 Gbit per second per pin-pair

- Two unidirectional point-to-point links
  - 2, 4, 8, 16, or 32 bits wide at 200 to 800Mhz (DDR) in each direction
  - Speed and Width are independently matched to upstream and downstream bandwidth needs (asymmetric)
  - Narrower devices can talk with wider devices
  - Width and speed negotiated at initialization (plug & play)
    - Link width equal to narrowest side of link established when exiting HyperTransport reset
    - Provision for using PnP BIOS to optimize both sides of link during configuration
- Data is packetized
  - Packets are multiples of 4-bytes in length
  - Serial interface with commands, addresses and data using the same wires

Scalable Bandwidth

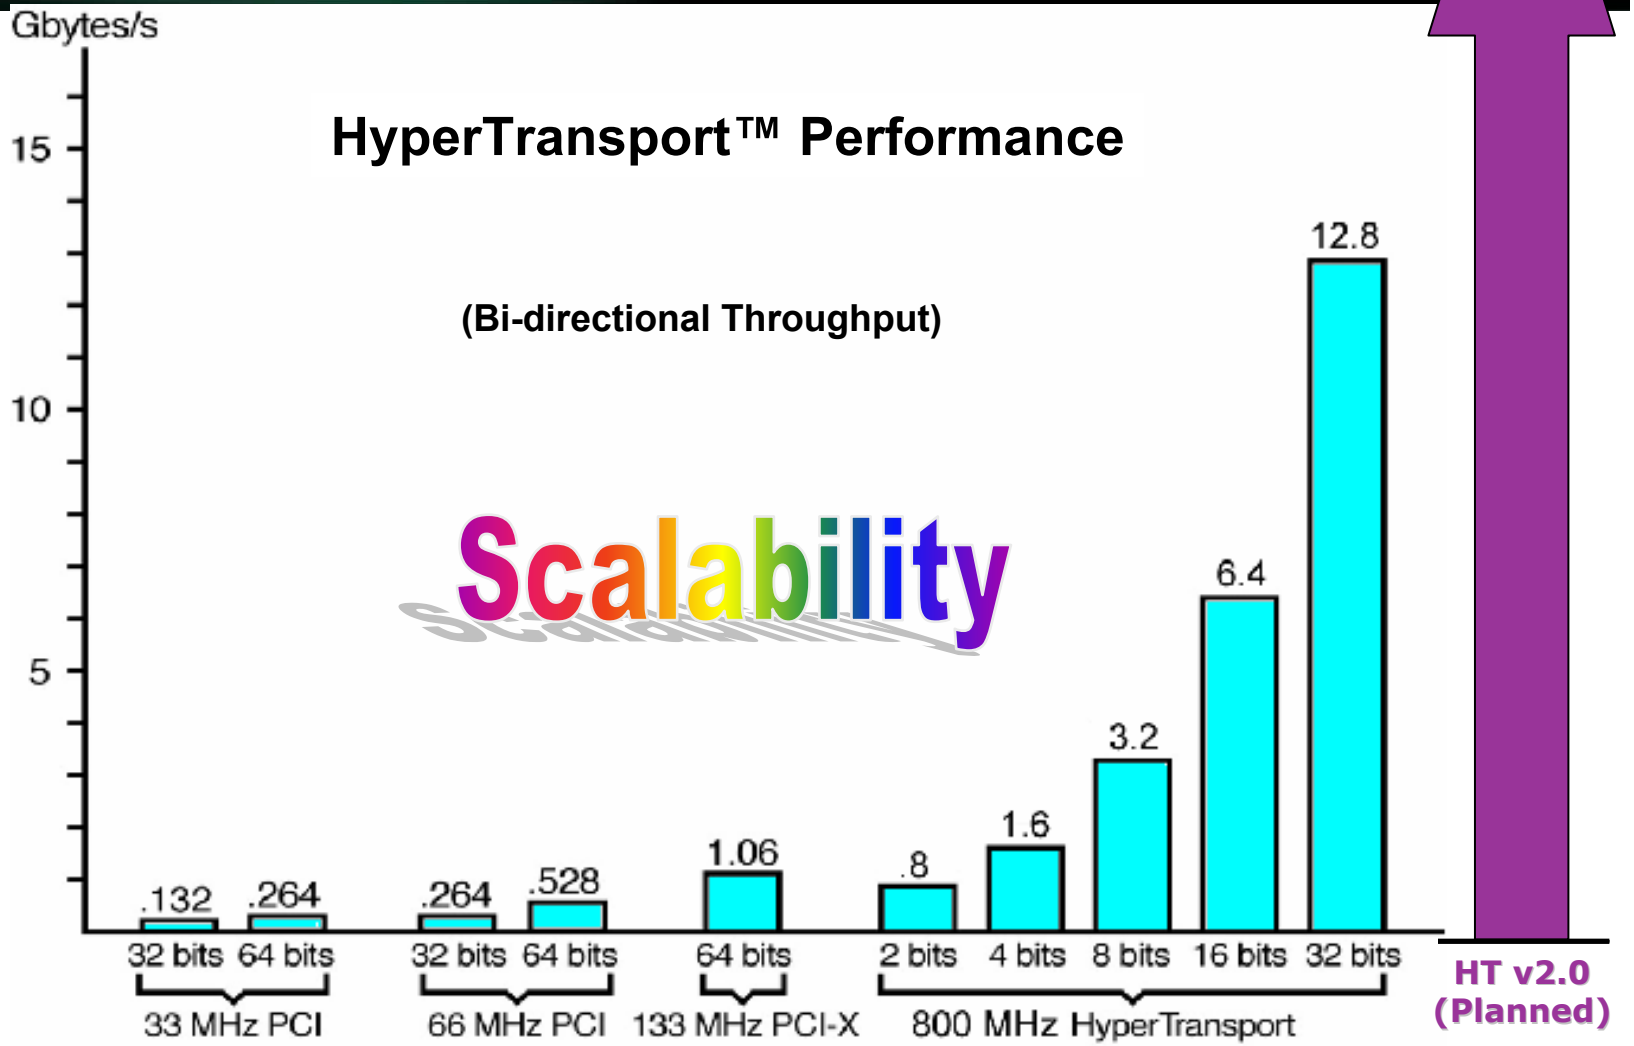

Clocks and Control Signals

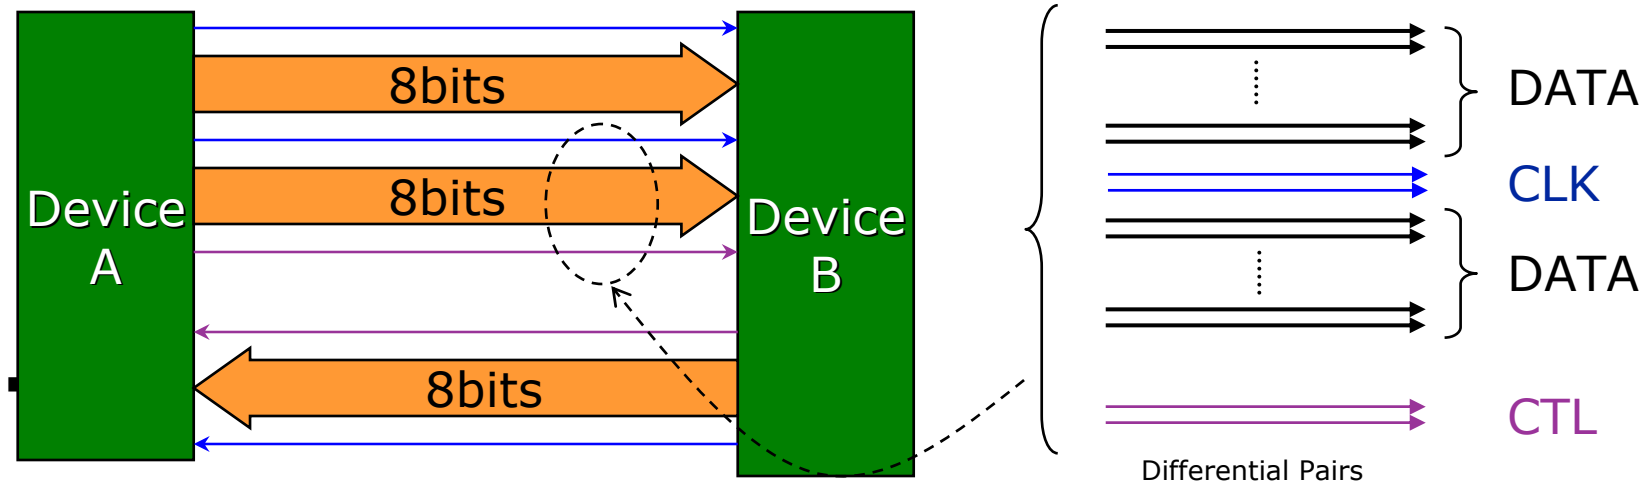

- Asynchronous clock forwarding
  - One clock is forwarded for each eight bits in each direction
  - Clocks are double pumped; an 800 MHz clock is used for 1.6Gbit/sec data rate
- Control line identifies command packets
  - De-asserted during data packets
- In-band system management & legacy signal transport
  - Eliminates sideband wires, interrupts use messages instead of wires
- Embedded code in back channel messages used for flow control
  - Code indicates how many buffers are available for each virtual channel

HyperTransport™ Pin Count

AMD DEVELOPER symposium 2002

- Additional HyperTransport™ signals
  - Power OK (PWROK)
  - Reset LDT (RESET_L)
- 55-pin HyperTransport bus provides 12x the bandwidth of PCI-32/33 with fewer pins!
- Signal to ground ratio is conservative at 4:1
- Optional link power down signals for mobile systems
  - LDTStop_L
  - DevReq_L
- Power per pin-pair is nil when in STOP state

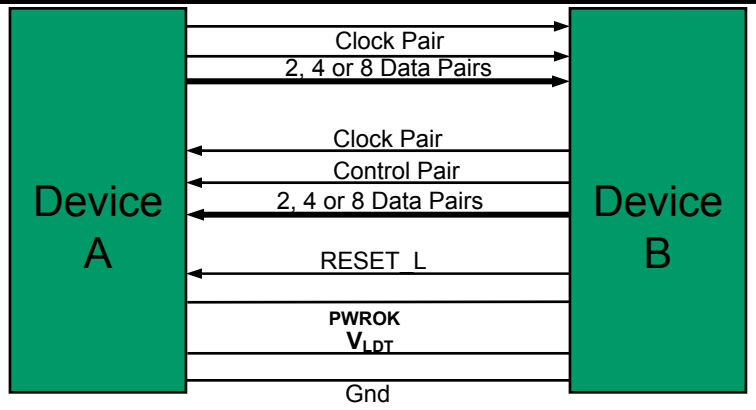

PWROK, RESET_L required for proper reset & init
V_{LDT} routed between devices is required for proper common mode range

Bus Width (Each Way)	2	4	8	16	32
Data Pins (total)	8	16	32	64	128
Clock Pins (total)	4	4	4	8	16
Control Pins (total)	4	4	4	4	4
Subtotal (high speed)	**16**	**24**	**40**	**76**	**148**
VLDT	2	2	3	6	10
GND	4	6	10	19	37
PWROK	1	1	1	1	1
RESET_L	1	1	1	1	1
Total Pins	**24**	**34**	**55**	**103**	**197**
Total Max BW GB/s	0.8	1.6	3.2	6.4	12.8

HyperTransport™ Tunnel

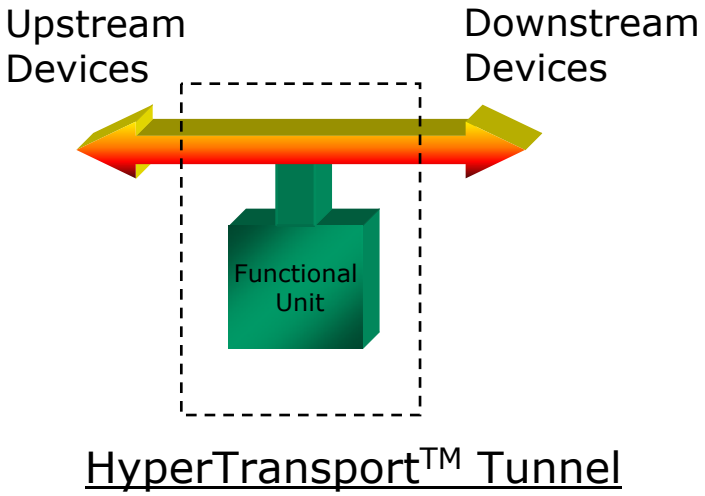

HyperTransport™ Tunnel

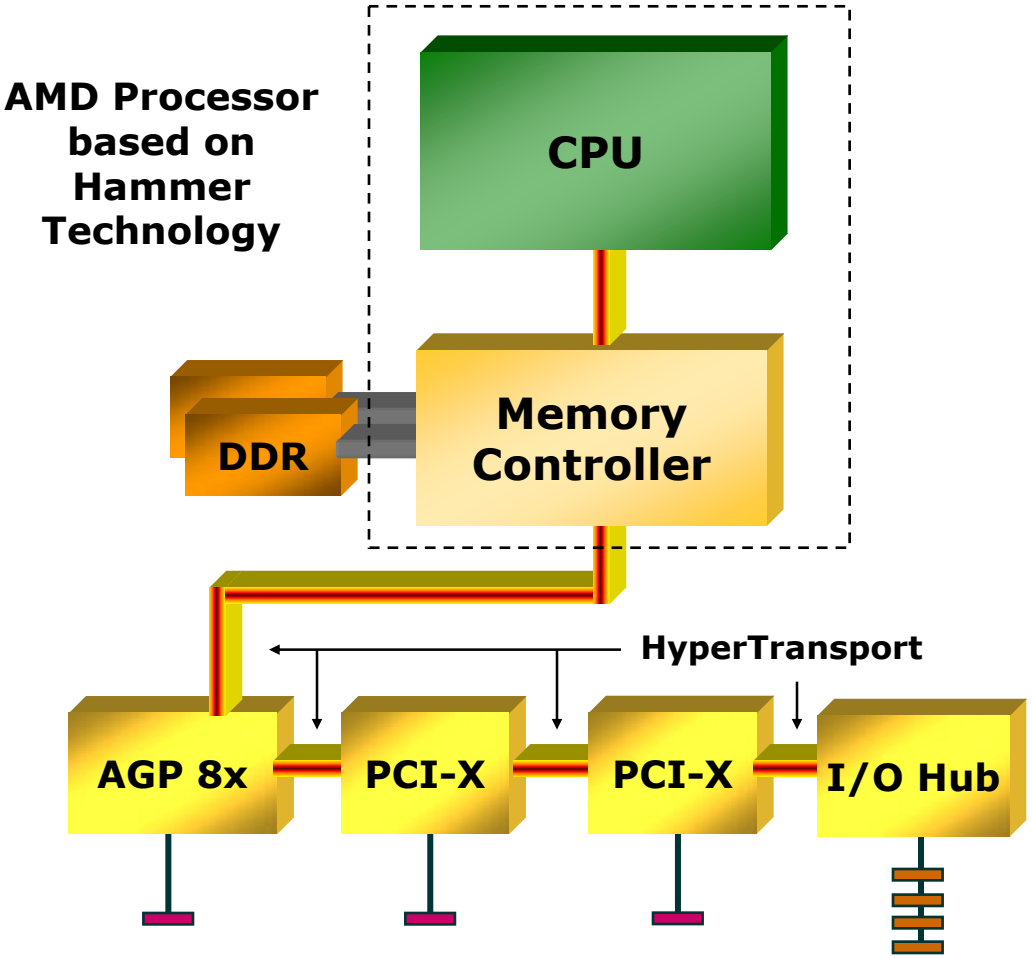

Tunneling Architecture

AMD DEVELOPER symposium 2002

- Enables a building block approach for system design
- Devices easily shared among different markets
  - Computation, Communication, Consumer…
- Unique "TUNNELING" capability provides easy expansion and simple board layout
- Multiple CPU and Memory Controller Architectures support same I/O component set
  - Multiple MIPS®, X86, etc…
- Larger component Total Available Market (TAM)
  - Lower cost components
  - Extended component production lifetime

Easy Board Layout

- Easily routed on 4 layer boards
- No special stack up considerations
- Minimal trace length issues
- Reduced skew considerations compared to current system buses
- Low noise and crosstalk

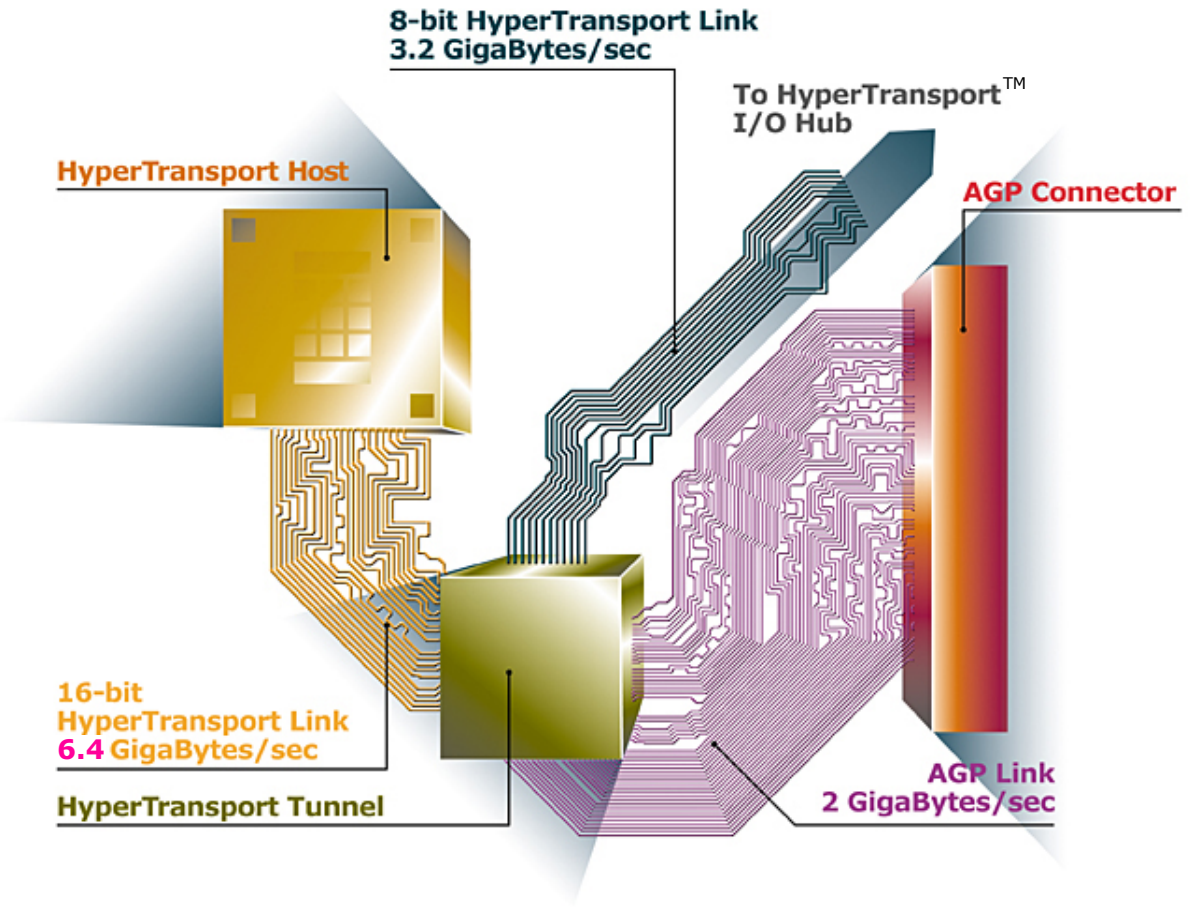

PCI Software Compatibility

- HyperTransport™ Technology builds upon PCI legacy
  - Supports standard Plug & Play enumeration
  - Transparent to Applications and Operating Systems
  - Minimal initialization required to optimize link efficiency
  - Supports PCI Power Management

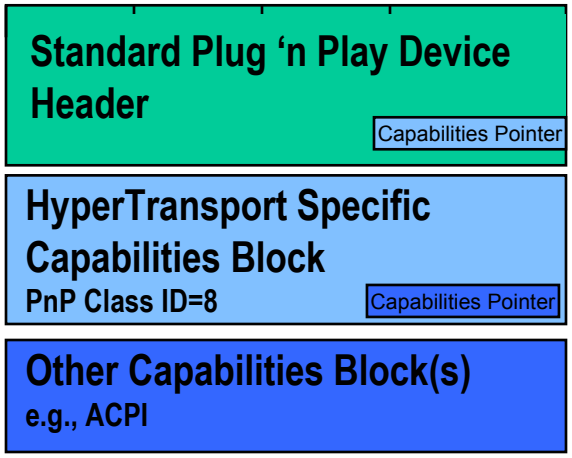

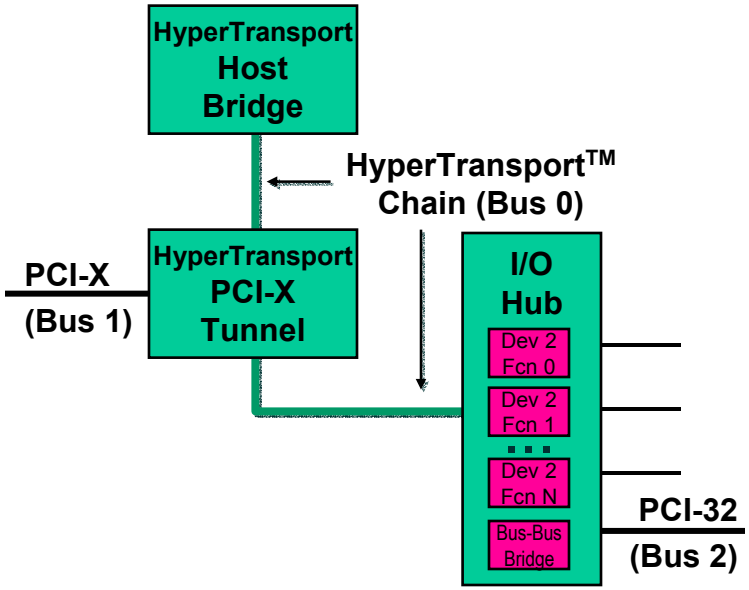

HyperTransport™ Data Structure

- Commands and interrupts are realized as a 32 bit command word
- Address and Data is preceded by a 64-bit header
  - 6-bit type field – Write, Read, Read Response, Fence & Flush
  - 26-bit Command specific field
  - 32-bit address field (command specific – Byte or DWORD)
- Data Blocks are 4 to 64 bytes long in multiple of 4 bytes
  - Transfers less than 4 bytes are padded

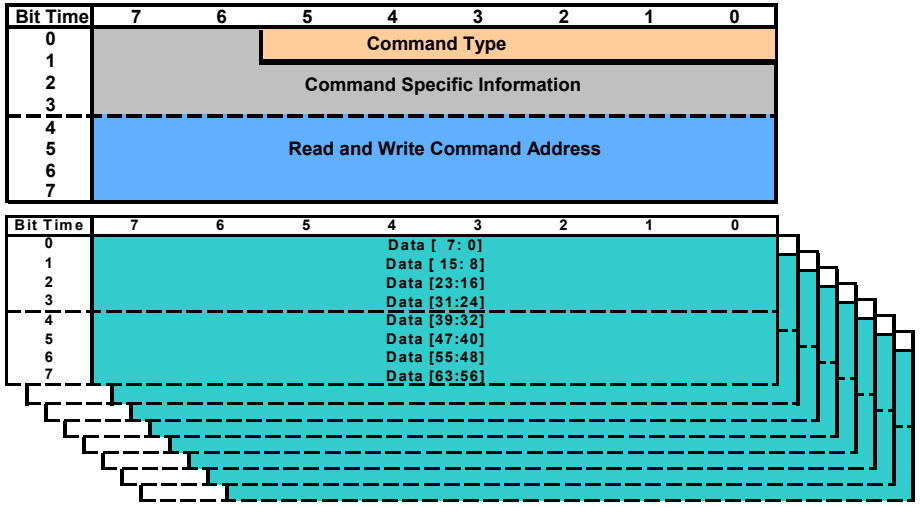

AMD DEVELOPER symposium 2002

HyperTransport™ Technology: Future Directions

Networking Requirements

AMD DEVELOPER symposium 2002

- Networks are growing in complexity and size
- Technologies like InfiniBand and 10Gbit Ethernet will provide increased bandwidth between devices (nodes) on these networks
- Network processors are becoming faster and more sophisticated
- Applications like security and streaming video are driving up the need for more bandwidth
- PCI is the common interconnect used today, but these new requirements are pushing the limits of its capabilities

HyperTransport™ technology provides an "in-the-box" solution while maintaining compatibility with PCI

HyperTransport™ Networking Extensions

AMD DEVELOPER symposium 2002

- NEW HyperTransport™ networking extensions improve performance for network applications
  - Control plane applications
  - Data plane applications
  - Look aside applications
- HyperTransport networking extensions enable one interface for all three applications, reducing complexity and increasing efficiency
- Networking extensions are designed to ensure interoperability with existing HyperTransport products
- Networking extensions planned to be available during the second half of 2002 as part of an update to the HyperTransport 1.x specification

HyperTransport™ Networking Extensions

AMD DEVELOPER symposium 2002

- Enhanced message passing protocol
  - Adds the ability to stream a sequence of packets to a given address
  - Number of streams limited only by device support (and 40 bit address space)
  - A set of messaging semantics reduce complexity and increase the efficiency of packet processing for communications applications
- Addition of 16 streaming point-to-point flow controlled virtual channels
  - Supports millions of end-to-end flow controlled individual streams on these 16 channels
  - Enables HyperTransport™ nodes to efficiently bridge SPI-4.2 traffic
  - Virtual channels can be assigned dedicated bandwidth with enhanced throughput for content-rich networking applications such as voice and video traffic
- Optional support for 64 bit addresses
  - HyperTransport v1.04 supports 40 bit addresses

HyperTransport™ Networking Extensions

AMD DEVELOPER symposium 2002

- Error handling protocol
  - Automatically detects and recovers from errors
  - Designed to increase the reliability and availability of the HyperTransport™ I/O connection
- Support for direct peer-to-peer transfers
  - Transfers can be sent directly between peers, without having to be reflected via the host, to increase efficiency
  - This enables high bandwidth look-up devices to cooperate directly with each other, reducing complexity
- Increased support for concurrent host transactions
  - Allows up to 1K transactions to be outstanding (current limit of 32)
  - Further improves flexibility and performance for communications applications
- Formalize the description of HyperTransport hubs and switches
  - Clearly defines the operation of HyperTransport switches and hubs within the standard and provides a clear roadmap for each application

AMD DEVELOPER symposium 2002

HyperTransport™ Technology System Design

1P AMD Athlon™ processor-based Desktop

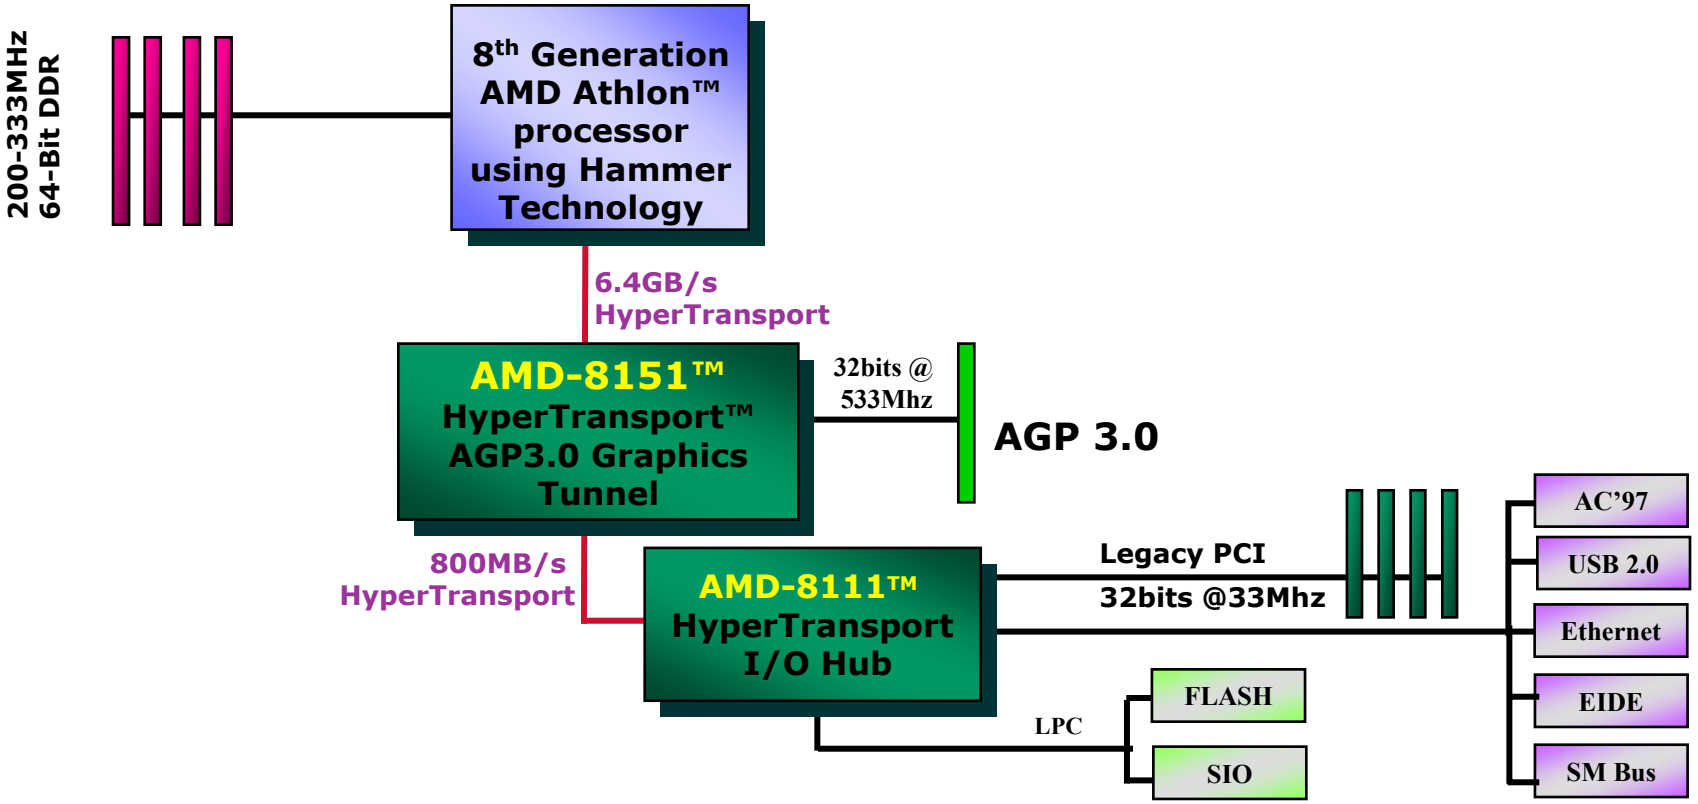

2P AMD Opteron™ processor-based Workstation

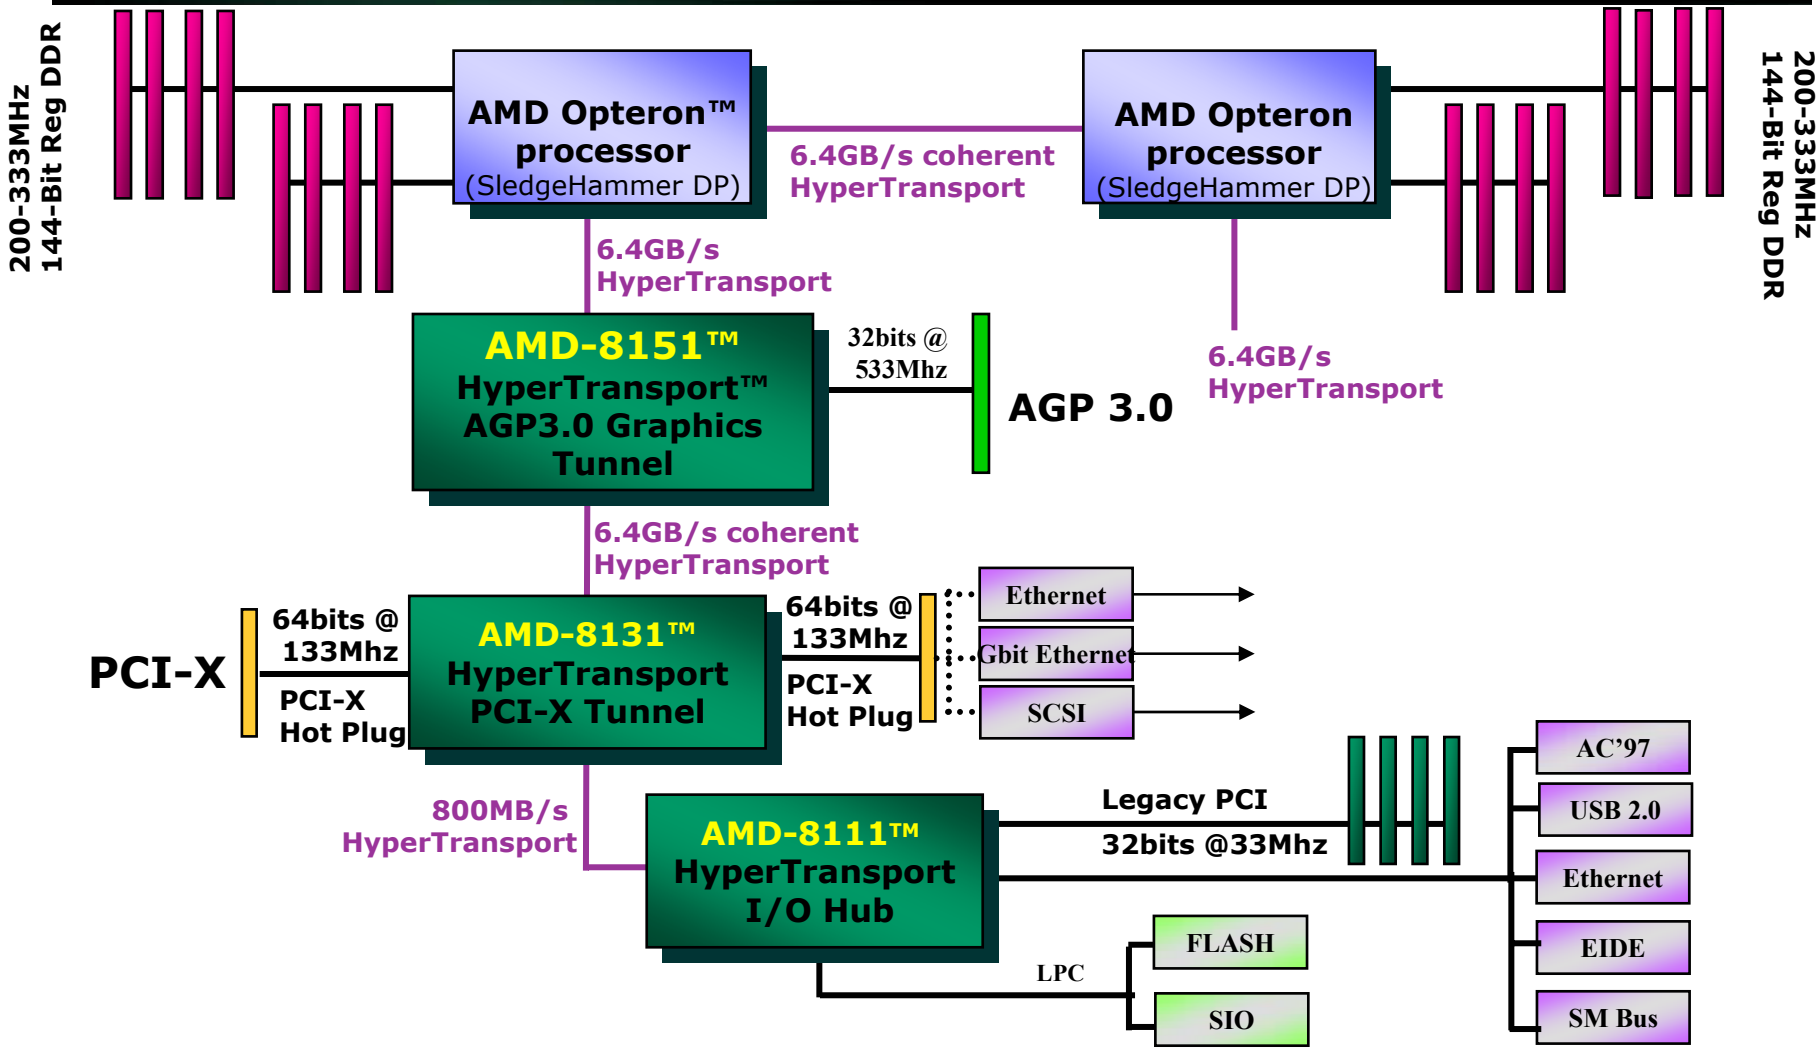

2P AMD Opteron™ processor-based Server

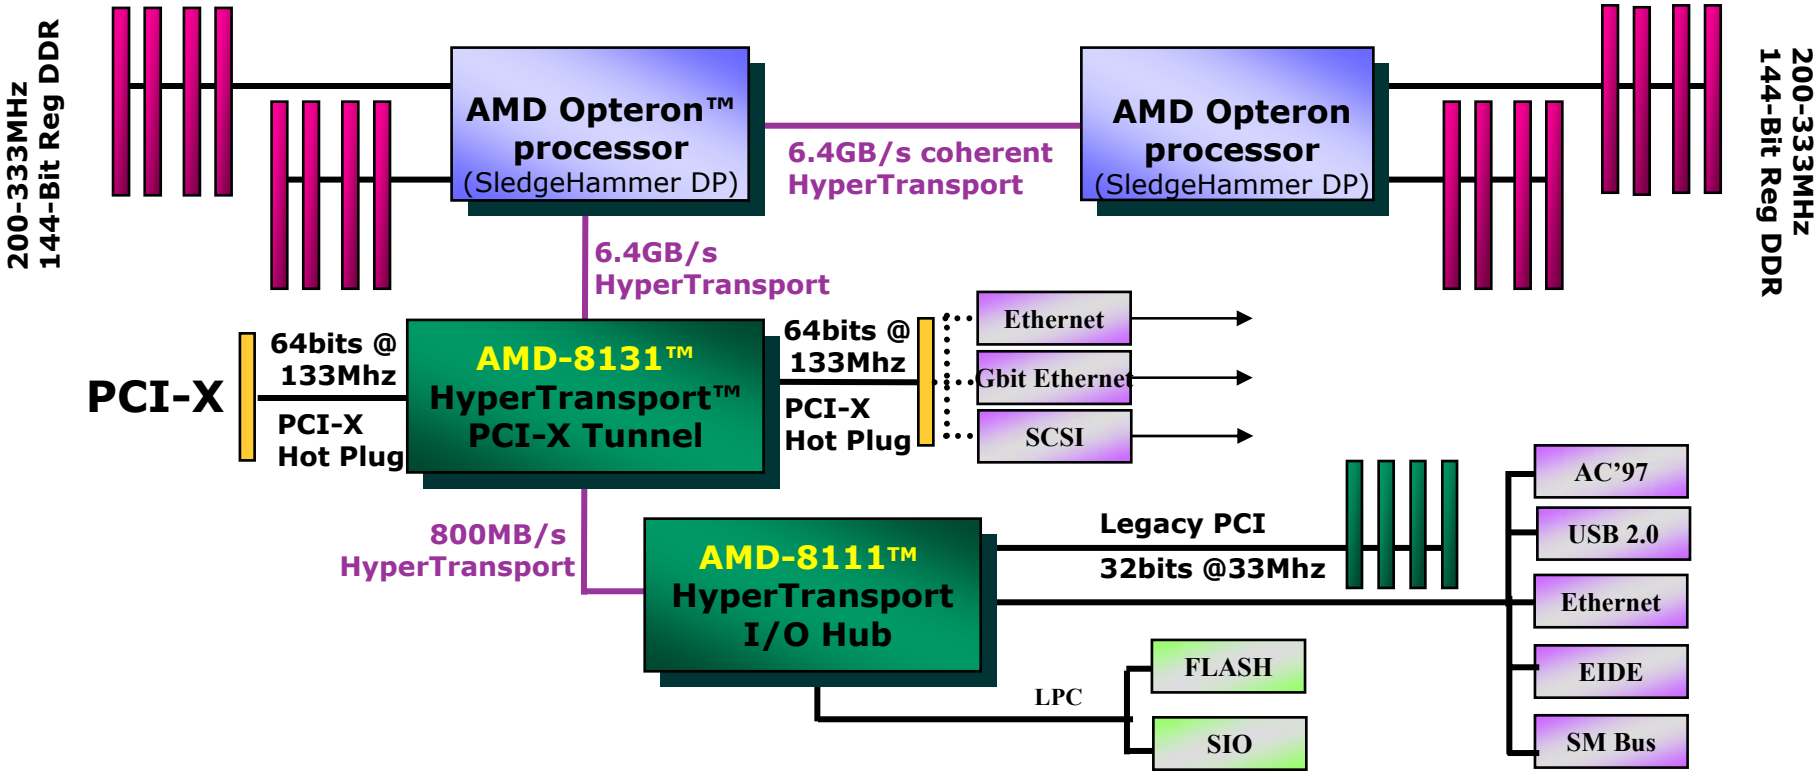

4P AMD Opteron™ processor-based Server

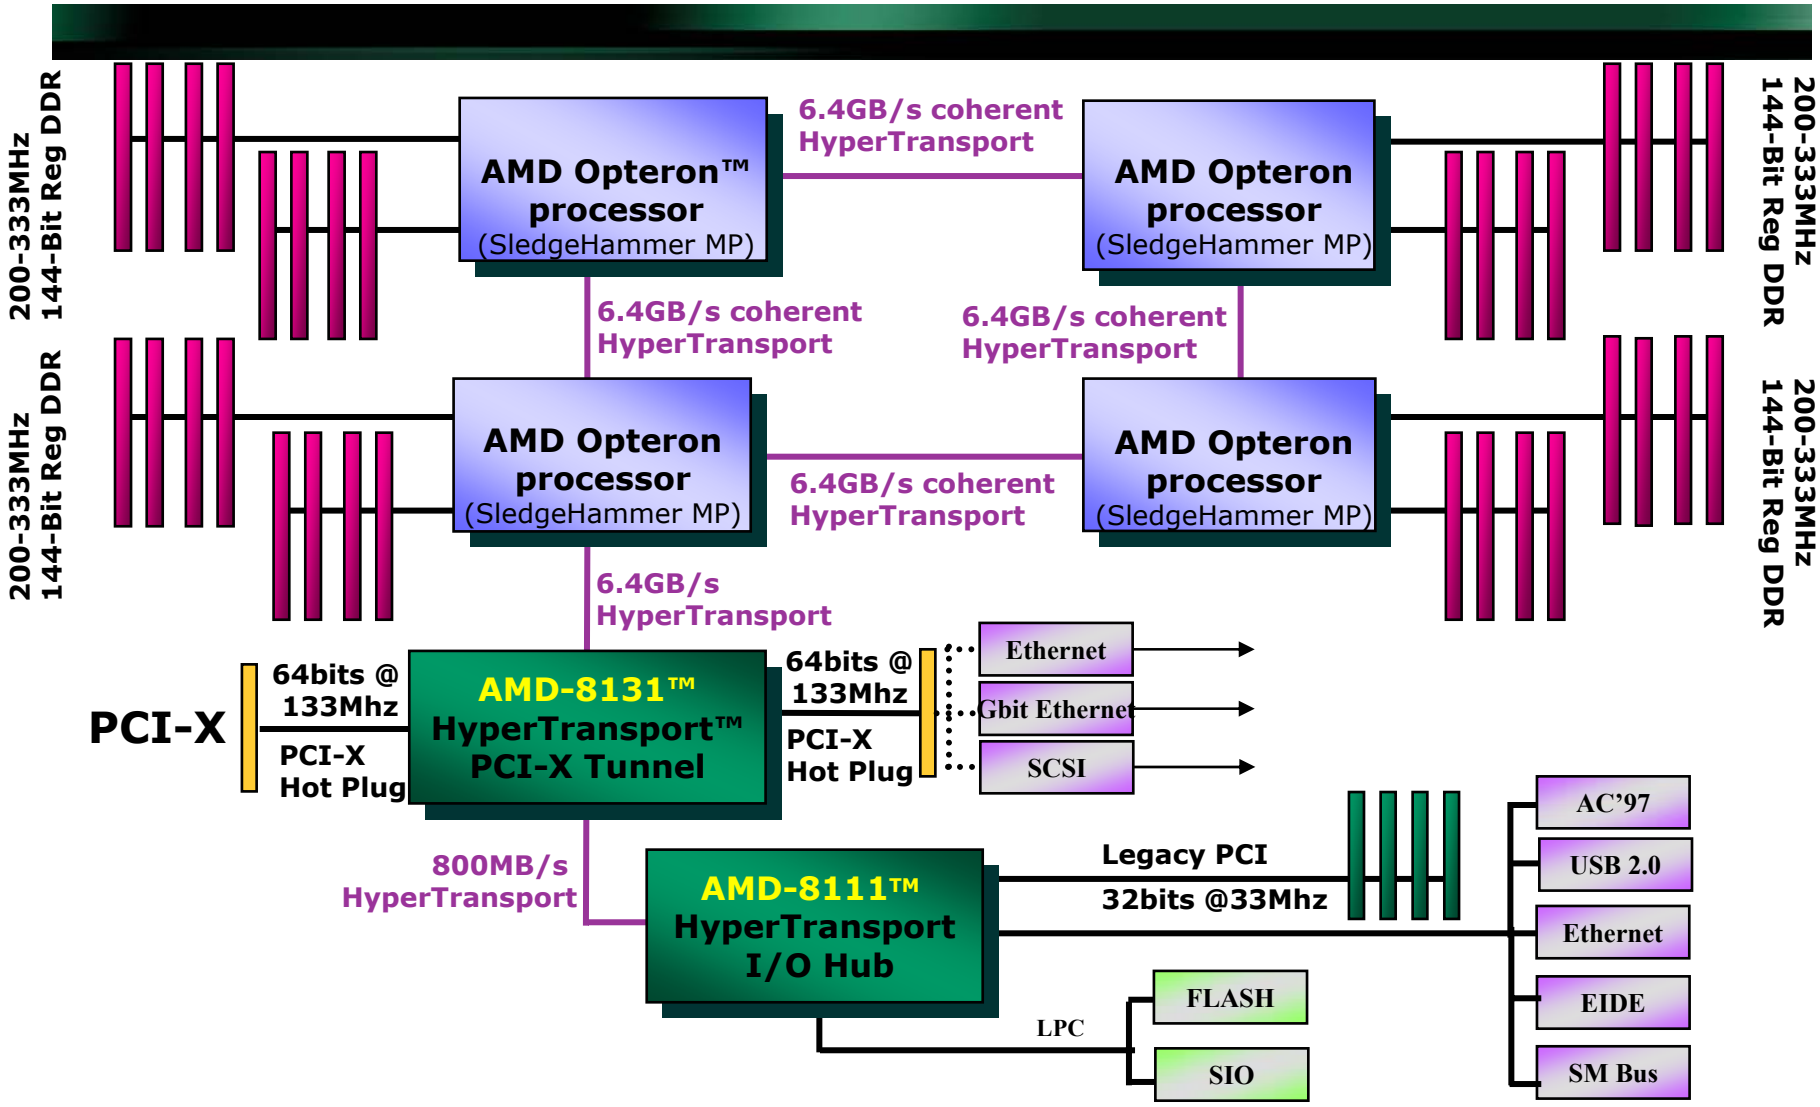

HyperTransport™ Technology Benefits

AMD DEVELOPER symposium 2002

- HyperTransport™ offers system designers
  - Flexibility
  - Scalability
  - Differentiation
  - New levels of performance
- HyperTransport offers chip designers
  - New markets
  - Large volumes
  - Increased interoperability
- HyperTransport offers end users
  - Choices
  - Reliability
  - Cost efficiency

Summary

AMD DEVELOPER symposium 2002

- HyperTransport™ Technology enables...
  - A significant increase in I/O bandwidth
  - A universal link that reduces the number of buses within the system
  - Support for HyperTransport tunnels that act as I/O building blocks
  - High performance bus for a wide range of applications
  - Highly scalable multiprocessing systems
- HyperTransport Technology is becoming a success story in the electronic industry
  - For the first time in the history of the Microprocessor, diverse CPU architectures are converging on a single bus
  - For the first time PC processors are going to have a packet-bus. This is changing the way people design computers

Trademark Attribution

AMD DEVELOPER symposium 2002